A Pilot mHealth Project for Monitoring Vital Body Signals and Skin Conditions

Rodrigue B. Tchema, Georgios Tzavellas, Marios Nestoros, and Anastasis C. Polycarpou(✉)

Department of Engineering, University of Nicosia, 1700 Nicosia, Cyprus
polycarpou.a@unic.ac.cy
http://www.unic.ac.cy

Abstract. In this paper, we present a pilot project on mobile health (mHealth) which is designed to monitor on a continuous basis the health condition of individuals using a mobile device such as a smartphone or a tablet. The objectives of this pilot project include highly accurate calculation of heart-beat rate using either a smartphone camera or an autonomous, self-powered mini-device which communicates measurement data to the mobile device through Bluetooth or Wi-Fi. In addition, we aim at detecting potentially dangerous skin conditions at an early stage using the smartphone camera and machine learning (ML) or deep learning (DL) algorithms. The trained algorithm will be able to detect malignant cases of skin conditions by searching through various built-in categories of commonly found skin disorders.

Keywords: Heart-beat rate detection · Skin disorder diagnosis · Deep learning algorithms

1 Introduction

As the human life expectancy is increasing, it is extremely important to limit heart diseases and other life-threatening health conditions, and help people maintain a healthy lifestyle. Unambiguously, the early detection of abnormalities in the health of individuals using existing technological tools will have a direct economic and social benefits. Mobile devices (like smartphones and tablets) have sensing elements (e.g., microphones, cameras, accelerometers, etc.) that can be effectively used to monitor vital signs of the human body. This sensor information can be further processed and communicated to the user or a healthcare professional through secured and fast communication channels, mobile applications, and cloud services allowing early intervention which can prevent health complications. The recent ubiquitous penetration of wearable sensors into the healthcare market, and their compatibility with mobile devices and software, allowed widespread use of mHealth monitoring and tracking.

The accuracy, sensitivity and reliability of these apps are often questionable by the majority of their users. Some wearables claim to provide a higher degree of

X. Gao et al. (Eds.): MobiHealth 2021, LNICST 440, pp. 270–280, 2022.
https://doi.org/10.1007/978-3-031-06368-8_18

accuracy in counting steps or pulse rate, and these come in the form of a watch, a wristband, or a strap around the chest. A recent study [1] on the accuracy of tracking apps and wearable devices has revealed that wearables provide a higher degree of accuracy with a maximum error rate equal to 5%. Apps, which depend on the smartphone's built-in accelerometer, are in general less accurate, and therefore, less reliable for professional use in sports.

The work described in this article concerns the use of a smartphone either as a gateway for transferring vital signals from a wearable or peripheral device to the cloud (for analysis and post-processing) or as a sensing device for recording and communicating health monitoring signals directly to the cloud server. At first, we present a prototype design of a compact and lightweight electronic circuit able to accurately extract the heart-beat rate based on the Photo-PlethysmoGraphy (PPG) technique. Then, we investigate the idea of using the smartphone's camera to capture a short video of the finger tip on top of the camera's lens, and use this information to obtain the pulse rate through suitable video processing and signal analysis. At last, we employ machine/deep learning algorithms on images taken using the smartphone's camera for early detection of skin malignancies such as cherry angioma and squamous cell carcinoma.

2 Project Activities

2.1 Heart-Beat Detection Using a Wireless Peripheral Mini-Device

For a daily monitor of someone's health, it is highly important that a reliable device is used to accurately calculate the heart-beat rate at a given moment. Not only that but the measurement data should be transferred to the cloud, post-processed, and compared with historical data thus providing mobile users important information on their health and well-being. Artificial intelligence (AI) algorithms could be employed in order to extract more meaningful results and predictions based on the overall picture of one's health.

In this project, we aim at designing a peripheral mini-device capable of detecting accurately the heart-beat rate of an individual under different body conditions (resting, sleeping, walking or running). For this purpose, we designed an electronic circuit to detect the weak pulse of the finger vein and convert that into a noise-free, amplified signal as a function of time. The technique we used is called Photo-PlethysmoGraphy (PPG) [2,3], which is based on the use of a monochromatic infrared (IR) light emitted by an IR-LED (Light Emitting Diode). This technique could be used in the reflection mode or the transmission mode (see Fig. 1). For example, in the transmission mode, an emitting diode is used on the front side of the finger and a photodiode or phototransistor is used on the back side as a detector. As the volume of the blood that flows through the veins is not constant as a function of time, the transmittance - at that particular wavelength - varies with time. This response is highly correlated with the heart beat. Similar observations are recorded for the reflection mode. Either approach is equally implemented in such applications. In our project, we used the reflection mode, therefore, both the emitting LED and phototransistor are positioned on the front side of the finger.

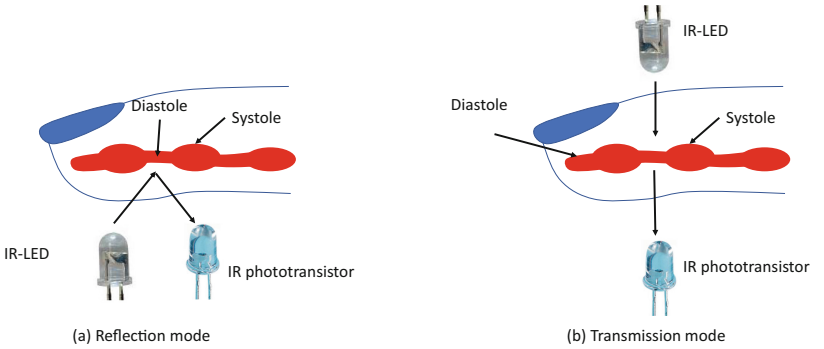

Fig. 1. Implementation of the Photo-PlethysmoGraphy (PPG) technique using the reflection and transmission modes.

A schematic diagram of the IR sensor electronic circuit augmented by active filtering and amplification is depicted in Fig. 2. The part on the left hand side corresponds to the IR sensor made of the IR-LED and the phototransistor. The wavelength emitted by the IR-LED is 940 nm. A band-pass filter (BPF) was designed to filter out unwanted noise and similar type of interference from sources other than the pulsating volume of blood in the finger. The filter was built to operate in a bandwidth of 40 bpm (beats per minute) to 200 bpm where 1 bpm corresponds to 1/60 Hz. In order to achieve a satisfactory level of filtering, two inverting active BPFs were used in conjunction with two inverting active amplifiers, as clearly shown in Fig. 2.

Circuit analysis of the proposed design results in an overall voltage amplification gain approximately equal to $A_v \approx 270$ and cut-off frequencies that define an operating bandwidth from 0.589 Hz to 3.386 Hz and a center frequency of 1.412 Hz. The frequency bandwidth corresponds to a heart-beat rate bandwidth between 35 bpm to 203 bpm with a peak rate at 85 bpm.

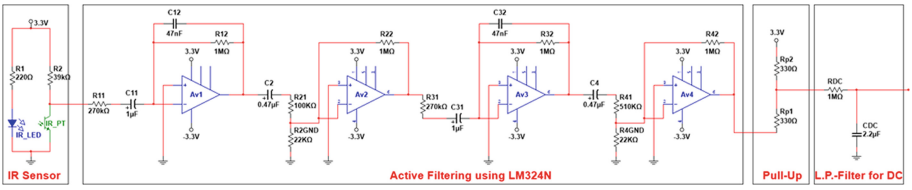

Fig. 2. Electronic circuit design of the heart-beat rate detection, filtering and amplification.

Revisiting the circuit in Fig. 2, two stages of active filtering and two stages of amplification were implemented in the final design. Initially, we started the design with a single active filter and a single amplification stage. However, it was

observed that the signal output, especially at low frequencies, had a significant interference from ambient light and other surrounding electromagnetic sources. In addition, the input signal from the IR sensor was too weak hence the need to improve the signal-to-noise ratio (SNR) was necessary. Therefore, two stages of filtering and amplification, using four concatenated Op-Amps, were implemented in order to resolve the aforementioned problem. In addition, two capacitors in between amplifiers and two grounded resistors were used in order to filter out the DC component and improve the SNR. The slight increase in power consumption due to the use of additional components was not considered problematic at this stage of the project.

The complete circuit, which was initially implemented on a breadboard, is shown in Fig. 3. As seen, there are two separate power supplies on the left bottom corner of the board. Specifically, there is a 5-V stabilizer that powers the microcontroller unit and the external SD-board along with the Real-Clock module, and one step-down power supply at 3.3 V that only powers the negative terminal of the quad Op-Amp. The positive terminal of electronic components such as filters, LEDs, and buzzers is attached to the 3.3 V step-down pin of the microcontroller board. The two-stage amplifier-filtering circuit appears in the center of the breadboard. Looking at the top left corner, there exist three optocouplers to be used at a later stage for switching on/off the Op-Amp, the 3.3 V step-down power supply, and the IR sensor. Although the SD-board and Real-Clock board are not essential for the key operations of the project, they are important for data storage in case there is loss of connection between the microcontroller and the Bluetooth receiver (e.g., computer, smartphone). In addition, LED indicators are used for possible debugging and component failure identification.

The circuit shown in Fig. 3 was tested in the lab and several PPG signals were obtained under different physical conditions. An example is shown in Fig. 4 clearly illustrating the systolic and diastolic phases of the heart pumping blood through the veins reaching all the way down to the finger tip. The signal shown in this figure corresponds to processed data which have been smoothed by software written on Python. Specifically, after the microcontroller sends the data sequence in a CSV text file, the software checks on the validity of the data and stores these values in memory. Then, with user's commands, the code implements smoothing and then plots the data as shown in Fig. 4. Data smoothing is based on the Savitzky-Golay filter library. This filter is implemented based on moving average and linear regression concepts. Two parameters are important to note. The first one is the window size that specifies the number of points used in the smoothing, whereas the second one is the degree of the polynomial implemented for curve fitting within the chosen window. In our case, the measurement window is 10 s, whereas the time it takes to transmit and save the data is approximately 5 s.

The heart-beat rate can be subsequently calculated by correctly identifying the extrema of the PPG signal. The maxima and minima of the signal can be easily calculated using a simple three-point algorithm. It simply checks the points before and after the current point. If their values are smaller than the current value, then the point is called a maximum. On the other hand, if their values

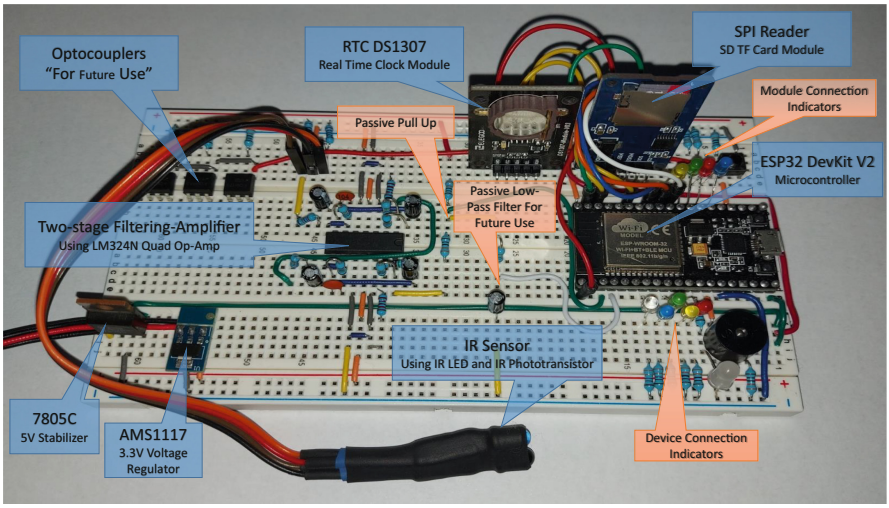

Fig. 3. Picture of the designed circuit implemented on a breadboard.

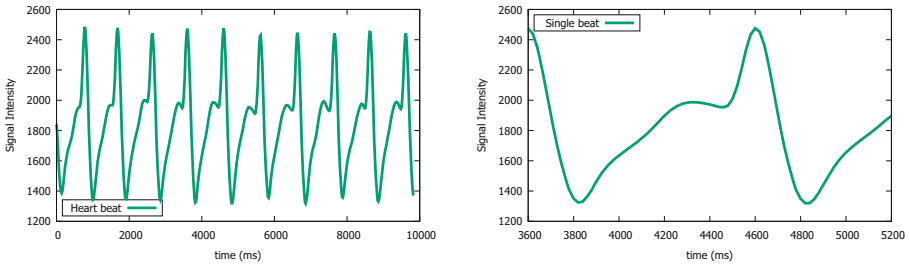

Fig. 4. Recorded heart beat pulses using the peripheral wireless mini-device: train of pulses (left); single pulse (right).

are larger than the current value, then the point is called a minimum. However, from the signal, it is easilly seen that we have additional local extrema. Thus, we save all extrema in two different arrays: one for local maxima and one for local minima. Then, the algorithm calculates the average value for each array and selects the values that are higher than the maxima average or lower than the minima average. Specifically, the heart-beat rate is calculated based on both maxima and minima, thus providing two estimates. If the deviation between the two estimates is larger than 2 bpm, then those estimates are discarded.

The accuracy of the prototype circuit was assessed by testing several users under resting and exercising conditions. Specifically, we tested six healthy individuals by recording their pulse rate when at rest and after performing minor exercise (10 or 20 push-ups). The recorded measurement was compared against the benchmark value obtained by pressing the index and middle fingers of the right hand on the opposite wrist and recording the number of beats manually in

60 s. As seen from Table 1, the two sets of data are in very close agreement. Specifically, the maximum percent deviation obtained was 3.61%, which is deemed to be satisfactory for the initial phase of the project. Further investigations will be performed with an increased number of subjects where the benchmark measurements will be recorded using an approved medical device.

Table 1. PPG circuit measurements and comparison with benchmark records.

User under test	Body activity*	Benchmark	Device	Percent deviation
User 'A'	Rest	67	66	1.52%
User 'A'	Ex-10	78	77	1.30%
User 'B'	Rest	72	71	1.41%
User 'B'	Ex-20	97	99	2.02%
User 'C'	Rest	64	62	3.23%
User 'C'	Ex-10	77	78	1.28%
User 'D'	Rest	86	73	3.61%
User 'E'	Rest	77	78	1.28%
User 'F'	Rest	73	72	1.39%

*Ex-10: After 10 push-ups; Ex-20: After 20 push-ups.

2.2 Heart-Beat Detection Using a Smartphone Built-In Camera

One of the main objectives of the project is to investigate the possibility of using the smartphone camera [4–6] to obtain accurate and reliable measurements of the pulse rate, as well as the oxygen concentration in blood. The smartphone camera is used to take a video of the finger tip which is directly attached to the lens. The flash light of the camera is always on during the recording time. The captured video files (.mp4 or .avi) from the camera will then be transferred to the cloud for further analysis and video processing, thereby allowing calculation of the heart-beat rate. At this point, our goal is to implement image processing techniques and dedicated algorithms that will provide accurate extraction of pulse rate first, and in the future, blood oxygen concentration. This is an alternative approach to calculating the pulse rate, as compared to the wireless circuit approach described in the previous section, aiming at providing a quick but accurate estimate of vital body signals when the user does not have access to the wearable peripheral device introduced earlier.

Our initial experimental data consists of approximately 40-seconds long videos of the index finger while in contact with the rear camera of the smartphone having the flashlight (LED) on. As light from the LED travels through tissue, it suffers absorption and scattering and its intensity is modulated by the pulsating flow of blood according to the PPG principle. The reflected light is

then recorded by the camera at 30 fps and the video is stored in the phone. The recording is then reduced into its frames and each frame is subsequently analysed into its three spectral components (Red, Green, Blue). For each component, we calculate the mean value of brightness per frame, and then, we apply a moving average filter in order to smooth the measurement data.

The next step is to identify the peaks in the signal and count their number in a given window. Knowing the number of peaks and the time span of the window, the pulse rate can be easily calculated. The peaks are identified by calculating the gradient (slope) at nearby points to the left and the right of the current evaluation point. This process takes place on data produced by the moving average filter, thus the data are smooth and free of high-frequency noise. A particular data point is identified as a peak (e.g., maximum) if the gradient is positive at three consecutive points before the peak and negative at three consecutive points following the peak. Taking three points before and after improves the calculation of the heart-beat rate by 5–10% as compared to taking a single point before and after. In our code, we experimented with a moving average based on the past 7 samples.

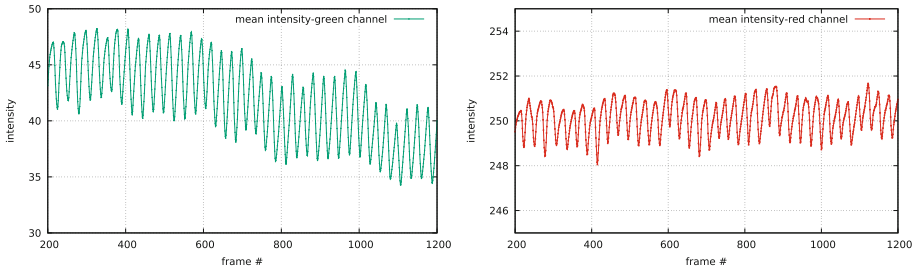

Fig. 5. Mean pixel intensity vs frame number as recorded at 30 fps from a rear smartphone camera: green channel (left); red channel (right). (Color figure online)

Once the video is recorded using the smartphone, the .mp4/.avi file is then transferred to a personal computer where all steps on video/image processing and data analysis are done using MATLAB codes. Figure 5 illustrates the GREEN and RED signal obtained from the smartphone camera after post-processing and smoothing. It is quite evident that both GREEN and RED signals present an oscillatory behavior as a result of the pulsating blood through the finger veins. As mentioned previously, the heart-beat rate can be calculated by counting the number of peaks within a given time window. At this point, it is worth noting that our near-term goal is to use available data from all three components (e.g., RED, GREEN, and BLUE) to also estimate oxygen concentration in blood.

The underlined method was tested against healthy people who belong to a wide range of age (17–52 years old). The recordings were performed under ambient light in the laboratory. The measurements are tabulated in Table 2. The results are compared against the pulse rate measured using the traditional

benchmark approach described earlier. Specifically, we used two types of smartphone: one based on iOS and another one based on ANDROID. From this exercise, it was observed that the percent deviation is relatively high in some cases reaching a value of up to 11.7%. Note that the BLUE component was excluded from the measurements because of its low correlation with pulse-rate calculation. Our observation is that the results are susceptible to slight movements of the finger, ambient light conditions, or even the charging level of the smartphone's battery. Consequently, it is highly important that, in the near future, all these effects be carefully filtered out from the individual RGB signals thus providing the ground for more accurate PR calculations.

Table 2. Pulse rate (PR) in bpm measured using a smartphone camera.

User	Device	PR - GREEN	PR - RED	Benchmark	Percent deviation
1	IOS	54	59	56	3.57% (G) : 5.36% (R)
2	IOS	64	65	66	3.03% (G) : 1.52% (R)
3	IOS	54	59	56	3.57% (G) : 5.36% (R)
4	IOS	86	83	83	3.61% (G) : 0.00% (R)
5	ANDROID	53	56	60	11.7% (G) : 6.67% (R)
6	ANDROID	62	62	67	7.46% (G) : 7.46% (R)
7	ANDROID	62	62	61	1.64% (G) : 1.64% (R)
8	ANDROID	62	60	65	4.62% (G) : 7.69% (R)

2.3 Deep Learning for Early Detection of Serious Skin Disorders

In this part of the project, we are using Deep Learning (DL) algorithms [7] to detect at an early stage possible health-threatening skin conditions in humans. We strongly believe that such an automated diagnosis system could be complementary to the work of experienced physicians, thereby increasing confidence in the results. In fact, there is enough evidence in the literature [8] that deep learning algorithms currently investigated in the area of medical diagnostics are equivalent in terms of performance to that of health-care professionals.

Health-threatening conditions may include categories of skin cancer some of which might be classified as malignant and life threatening. The smartphone camera is used to take a picture of the affected skin area, which is then uploaded to the cloud for processing and classification. The objective of this work is to implement codes based on DL algorithms able to identify with high success rate the type of skin condition. In this effort, we implemented YOLO v.4 (You Only Look Once) which is considered as the state-of-art algorithm in object detection due to its high accuracy and superb speed. The methodology used in

implementing this particular algorithm include the following: (a) Data collection; (b) Data processing which involves sorting, labeling, and dataset augmentation; (c) Classification.

Our first task was to collect images from different repositories available to the public or researchers. Initially, we chose to work with only four classes of skin conditions namely angioma cherry, squamous cell carcinoma, varicella, and normal skin. Once we collected a number of images from different sites, we started the pre-processing of images from which we excluded repeated ones and those with very bad resolution. Eventually, we ended up with a balanced dataset containing 120 images per class (480 images for the four classes in total). At this stage, we performed four simulations, each one with a different number of training sets. In other words, the training dataset was constantly modified but the testing dataset was kept fixed to a set of 60 images for all simulations. In order to test our approach in terms of performance accuracy, we restricted the number of classes to just two (angioma cherry and squamous cell carcinoma). For data augmentation, we used a MATLAB code. Particularly, images were flipped, rotated and cropped. For annotation, we used an open-source software called *LabelImg* which has the option of saving data into YOLO format.

For the first simulation, we implemented data augmentation on the 120 images per class, thus raising the total number of images per class to 200. For the two classes in hand, we had 400 images to use for training. In order to annotate the images, we used a bounding box around the most pronounced lesion in a given image, as shown in Fig. 6. Upon the completion of the training process, a total of 60 images per class were used for the testing phase. The result was 96.7% accuracy for the case of angioma cherry and 10% accuracy for the case of squamous cell carcinoma. Specifically, out of 60 images for each class, 58 images were correctly classified as angioma cherry and only 6 images for squamous cell carcinoma. This indicates very good performance on the classification of angioma cherry but very poor performance on the classification of squamous cell carcinoma.

In an attempt to improve the above results, a second simulation was performed. As the YOLO algorithm is used for classification instead of detection, the labeling technique was altered. Instead of labeling only the more pronounced lesions on an image (as shown in Fig. 6), we chose to label the entire image. In addition, a third class of training data was added which is the normal skin (image without any abnormality). A total of 120 images for all three classes was initially trained. The corresponding results in terms of success rate include 76.7% for angioma cherry, 53.3% for squamous cell carcinoma, and 100% for normal skin. For further improvement of the results, we repeated the simulation with a training dataset of 360 images (120 images per class). In this case, the achieved success rate was 90.0% for angioma cherry, 73.3% for squamous cell carcinoma, and 100% for normal skin. At last, the dataset was increased to a total of 1800 images (600 images per class). This increase in the training set resulted in a significant improvement in the success rate. Specifically, the obtained success rate for angioma cherry was 98.3%, for squamous cell carcinoma was 76.0%, and for

Fig. 6. Assigning bounding box around most pronounced lesion in an image: angioma cherry (left); Squamous cell carcinoma (right).

normal skin was 100%. The increase in the number of images per simulation was achieved by flipping an image, rotating an image, or by adding Gaussian noise. Consequently, we generate additional images from the initial group of images obtained from available repositories. The results are collectively tabulated in Table 3.

Table 3. Success rate of classification based on different size of training sets.

Size of training set	Angioma cherry	Squamous cell carcinoma	Normal skin
120	76.7%	53.3%	100%
360	90.0%	73.3%	100%
1800	98.3%	76.0%	100%

3 Conclusions

In this paper, we presented preliminary results obtained in the context of a recently launched pilot project on mobile health. The results sought are quite promising providing a solid ground for further investigation and improvement of the techniques utilized in this work. As already indicated, our primary goal is to enhance currently used techniques in order to provide accurate estimates of vital health signals or even to use newly introduced methods, such as machine

learning or similar algorithms, in order to increase the probability of accurately identifying harmful skin conditions for the human being. As illustrated in this paper, a newly proposed electronic circuit used as a peripheral device was developed illustrating accurate extraction of the heart beat rate. This circuit will be augmented in the near term to calculate blood oxygen saturation as well. We also demonstrated a MATLAB algorithm which provides accurate calculation of heart beat rate based on image processing of a video of one's finger directly attached to the lens of a smartphone camera. This algorithm will be extended soon to also calculate the blood oxygen saturation. It was also illustrated that machine learning algorithms could be used as a tool for potential diagnosis of harmful skin conditions at an early stage. For a successful implementation of these algorithms and improved outcome, it is highly important that data are judiciously prepared based on certain rules and criteria before they are eventually fed to these algorithms for training and testing.

Acknowledgments. The authors would like to thank the University of Nicosia for their financial support of the project.

References

1. Stavropoulos, T.G., Andreadis, S., Mpaltadoros, L., Nikolopoulos, S., Kompatsiaris, I.: Wearable sensors and smartphone apps as pedometers in eHealth: a comparative accuracy, reliability and user evaluation. In: 2020 IEEE International Conference on Human-Machine Systems (ICHMS), Rome, pp. 1–6 (2020)
2. Daimiwal, N., Sundhararajan, M., Shriram, R.: Respiratory rate, heart rate and continuous measurement of BP using PPG. In: International Conference on Communication and Signal Processing, Melmaruvathur, India, pp. 999–1002 (2014)
3. Fujita, D., Suzuki, A.: Evaluation of the possible use of PPG waveform features measured at low sampling rate. IEEE Access **7**, 58361–58367 (2019)
4. Grimaldi, D., Kurylyak, Y., Lamonaca, F., Nastro, A.: Photoplethysmography detection by smartphone's videocamera. In: 6th IEEE International Conference on Intelligent Data Acquisition and Advanced Computing Systems, Prague, pp. 15–17 (2011)
5. Holz, C., Ofek, E.: Doubling the signal quality of smartphone camera pulse oximetry using the display screen as a controllable selective light source. In: 40th Annual International Conference of the IEEE Engineering in Medicine and Biology, Honolulu, pp. 1–4 (2018)
6. Pelegris, P., Banitsas, K., Orbach, T., Marias, K.: A novel method to detect heart beat rate using a mobile phone. In: 32nd Annual International Conference of the IEEE EMBS Buenos Aires, Argentina, pp. 5488–5491 (2010)
7. Winkler, J.K., et al.: Melanoma recognition by a deep learning convolutional neural network performance in different melanoma subtypes and localisations. Eur. J. Cancer **127**, 21–29 (2020)
8. Liu, X., et al.: A comparison of deep learning performance against health-care professionals in detecting diseases from medical imaging: a systematic review and meta-analysis. Lancet Digit. Health **1**(6), e271–e297 (2019)